

WHAT IS CLAIMED IS:

1. A mass storage controller system, comprising:
 2. a plurality of controllers for controlling an array of storage devices, each of
 3. the plurality of controllers comprising:
 4. a CPU for controlling the operation of a controller;
 5. program memory, coupled to the CPU, for storing program instructions
 6. and variables for the operation of the CPU; and
 7. cache memory, coupled to the CPU, for storing information related to
 8. the array of storage devices;
 9. wherein a controller of the plurality of controllers initiates a task to be
 10. performed, the controller initiating the task establishes a task coordination data
 11. object shared by the plurality of controllers, wherein the task coordination data
 12. object represents discrete partitions of the task to be performed and states for each
 13. partition, and wherein a free controller of the plurality of controllers selects a partition
 14. of the task available for processing as indicated by the states.
1. 2. The mass storage controller system of claim 1, wherein the state
2. indicates whether a partition is READY, IN PROGRESS, or COMPLETE.
1. 3. The mass storage controller system of claim 2, wherein a controller
2. selects a partition by examining the partitions in a READY state and selecting at
3. least one partition in the READY state to operate on.

1 4. The mass storage controller system of claim 3, wherein a partition is in
2 an IN PROGRESS state during processing.

1 5. The mass storage controller system of claim 4, wherein a controller
2 sets the partition selected for processing to a COMPLETE state upon completion of
3 processing for a partition.

1 6. The mass storage controller system of claim 1, wherein a controller
2 selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 7. The mass storage controller system of claim 1, wherein a partition is in
2 an IN PROGRESS state during processing.

1 8. The mass storage controller system of claim 1, wherein a controller
2 sets the partition selected for processing to a COMPLETE state upon completion of
3 processing for a partition.

1 9. The mass storage controller system of claim 1, wherein the states
2 provide a semaphore-mechanism for allowing a controller to ascertain whether to
3 acquire control over a partition.

1 10. The mass storage controller system of claim 1, wherein the initiating
2 controller is notified when all partition states are COMPLETE and performs whatever
3 completion actions are required.

1 11. The mass storage controller system of claim 1, wherein the task
2 coordination data object includes information about an operation to be performed
3 and a data set to be operated on.

1 12. A mass storage array subsystem, comprising:
2 a plurality of storage devices;
3 a backplane adapted to couple to said plurality of storage devices; and
4 a plurality of controllers, coupled to the backplane, for controlling the plurality
5 of storage devices, the plurality of controllers having a first interface to couple to a
6 host system and a second interface adapted to couple to said backplane to
7 communicate with said plurality of storage devices;
8 wherein each of the plurality of controllers comprise a CPU for controlling the
9 operation of a controller, program memory for storing program instructions and
10 variables for the operation of the CPU and cache memory for storing information
11 related to the array of storage devices, and wherein a controller of the plurality of
12 controllers initiates a task to be performed, the controller initiating the task
13 establishes a task coordination data object shared by the plurality of controllers,
14 wherein the task coordination data object represents discrete partitions of the task to
15 be performed and states for each partition, and wherein a free controller of the
16 plurality of controllers selects a partition of the task available for processing as
17 indicated by the states.

1 13. The mass storage array subsystem of claim 12, wherein the state
2 indicates whether a partition is READY, IN PROGRESS, or COMPLETE.

1 14. The mass storage array subsystem of claim 13, wherein a controller
2 selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 15. The mass storage array subsystem of claim 14, wherein a partition is
2 in an IN PROGRESS state during processing.

1 16. The mass storage array subsystem of claim 15, wherein a controller
2 sets the partition selected for processing to a COMPLETE state upon completion of
3 processing for a partition.

1 17. The mass storage array subsystem of claim 12, wherein a controller
2 selects a partition by examining the partitions in a READY state and selecting at
3 least one partition in the READY state to operate on.

1 18. The mass storage array subsystem of claim 12, wherein a partition is
2 in an IN PROGRESS state during processing.

1 19. The mass storage array subsystem of claim 12, wherein a controller
2 sets the partition selected for processing to a COMPLETE state upon completion of
3 processing for a partition.

1 20. The mass storage array subsystem of claim 12, wherein the states
2 provide a semaphore-mechanism for allowing a controller to ascertain whether to
3 acquire control over a partition.

1 21. The mass storage array subsystem of claim 12, wherein the initiating
2 controller is notified when all partition states are COMPLETE and performs whatever
3 completion actions are required.

1 22. The mass storage array subsystem of claim 12, wherein the task
2 coordination data object includes information about an operation to be performed
3 and a data set to be operated on.

1 23. A method for cooperative distributed task management in a storage
2 subsystem with multiple controllers using cache locking, comprising:
3 initiating by an initiating controller a task to be performed;
4 establishing by the initiating controller a task coordination data object shared
5 by the multiple controllers, wherein the task coordination data object represents
6 discrete partitions of the task to be performed and states for each partition; and
7 selecting by a free controller a partition of a task available for processing as
8 indicated by the states.

1 24. The method of claim 23 further comprising indicating a state of a
2 partition as being READY, IN PROGRESS, or COMPLETE.

1 25. The method of claim 24 wherein the selecting by a free controller is
2 performed by examining the partitions in a READY state and selecting at least one
3 partition in the READY state to operate on.

1 26. The method of claim 25, wherein a partition is in an IN PROGRESS
2 state during processing.

1 27. The method of claim 26 further comprising setting by a controller a
2 partition selected for processing to a COMPLETE state upon completion of
3 processing for the partition.

1 28. The method of claim 23, wherein the selecting by a free controller is
2 performed by examining the partitions in a READY state and selecting at least one
3 partition in the READY state to operate on.

1 29. The method of claim 23, wherein a partition is in an IN PROGRESS
2 state during processing.

1 30. The method of claim 23 further comprising setting by a controller a
2 partition selected for processing to a COMPLETE state upon completion of
3 processing for the partition.

1 31. The method of claim 23 wherein the states provide a semaphore-
2 mechanism for allowing a controller to ascertain whether to acquire control over a
3 partition.

1 32. The method of claim 23 further comprising notifying the initiating
2 controller when all partition states are complete and performing completion actions
3 that are required.

1 33. The method of claim 23, wherein the task coordination data object
2 includes information about an operation to be performed and a data set to be
3 operated on.

1 34. An article of manufacture comprising a program storage medium
2 readable by a computer, the medium tangibly embodying one or more programs of
3 instructions executable by the computer to perform a method for cooperative
4 distributed task management in a storage subsystem with multiple controllers using
5 cache locking, the method comprising:
6 initiating by an initiating controller a task to be performed;
7 establishing by the initiating controller a task coordination data object shared
8 by the multiple controllers, wherein the task coordination data object represents
9 discrete partitions of the task to be performed and states for each partition; and
10 selecting by a free controller a partition of a task available for processing as
11 indicated by the states.

1 35. The article of manufacture of claim 34 further comprising indicating a
2 state of a partition as being READY, IN PROGRESS, or COMPLETE.

1 36. The article of manufacture of claim 35 wherein the selecting by a free
2 controller is performed by examining the partitions in a READY state and selecting
3 at least one partition in the READY state to operate on.

1 37. The article of manufacture of claim 36, wherein a partition is in an IN
2 PROGRESS state during processing.

1 38. The article of manufacture of claim 37 further comprising setting by a
2 controller a partition selected for processing to a COMPLETE state upon completion
3 of processing for the partition.

1 39. The article of manufacture of claim 34, wherein a free controller
2 selects a partition by examining the partitions in a READY state and selecting at
3 least one to operate on.

1 40. The article of manufacture of claim 34 further comprising notifying the
2 initiating controller is notified when all partition states are complete and performs
3 whatever completion actions required.

1 41. The article of manufacture of claim 34, wherein the task coordination
2 data object includes information about an operation to be performed and a data set
3 to be operated on.

1 42. A mass storage controller system, comprising:
2 a plurality of intermediate controller means for controlling an array of storage
3 devices, each of the plurality of controllers comprising:
4 CPU means for controlling the operation of a controller;
5 program memory means, coupled to the CPU means, for storing
6 program instructions and variables for the operation of the CPU; and
7 cache memory means, coupled to the CPU means, for storing
8 information related to the array of storage devices;
9 wherein an intermediate controller means of the plurality of intermediate
10 controller means initiates a task to be performed, the intermediate controller means
11 initiating the task establishes a task coordination data object shared by the plurality
12 of intermediate controller means controller means, wherein the task coordination
13 data object represents discrete partitions of the task to be performed and states for
14 each partition, and wherein a free intermediate controller means selects a partition
15 of the task available for processing as indicated by the states.